

A few remarks on DOE and Computing

Simulation is at the core of our mission success

- *Many of the mission areas, from energy and emergency response/preparedness, to nuclear security and science, involve classes of problems that cannot be tested other than virtually.*
- *DOE and the labs have equities in outcome; Context of the missions provides intellectual headroom and career long commitments to the outcomes*
- *Go to agency for informing urgent decisions*

This is the only way to approach problems that cannot be instrumented

- *Many of the mission areas, from energy and emergency response/preparedness, to nuclear security and science; classes of problems that cannot be tested other than virtually.*
- *Vital to labs. Built around Uncertainty Quantification, Verification & Validation*
- *Many successes. Billions in deferred costs. Likely more in the future*
- *Innovation in our missions will require simple tools to explore ideas*
- *Rigor in prediction. A place we must continue to pioneer.*
- *There is a cost for inaction.*

Sector-specific agency for Energy Infrastructure (PDD-21)

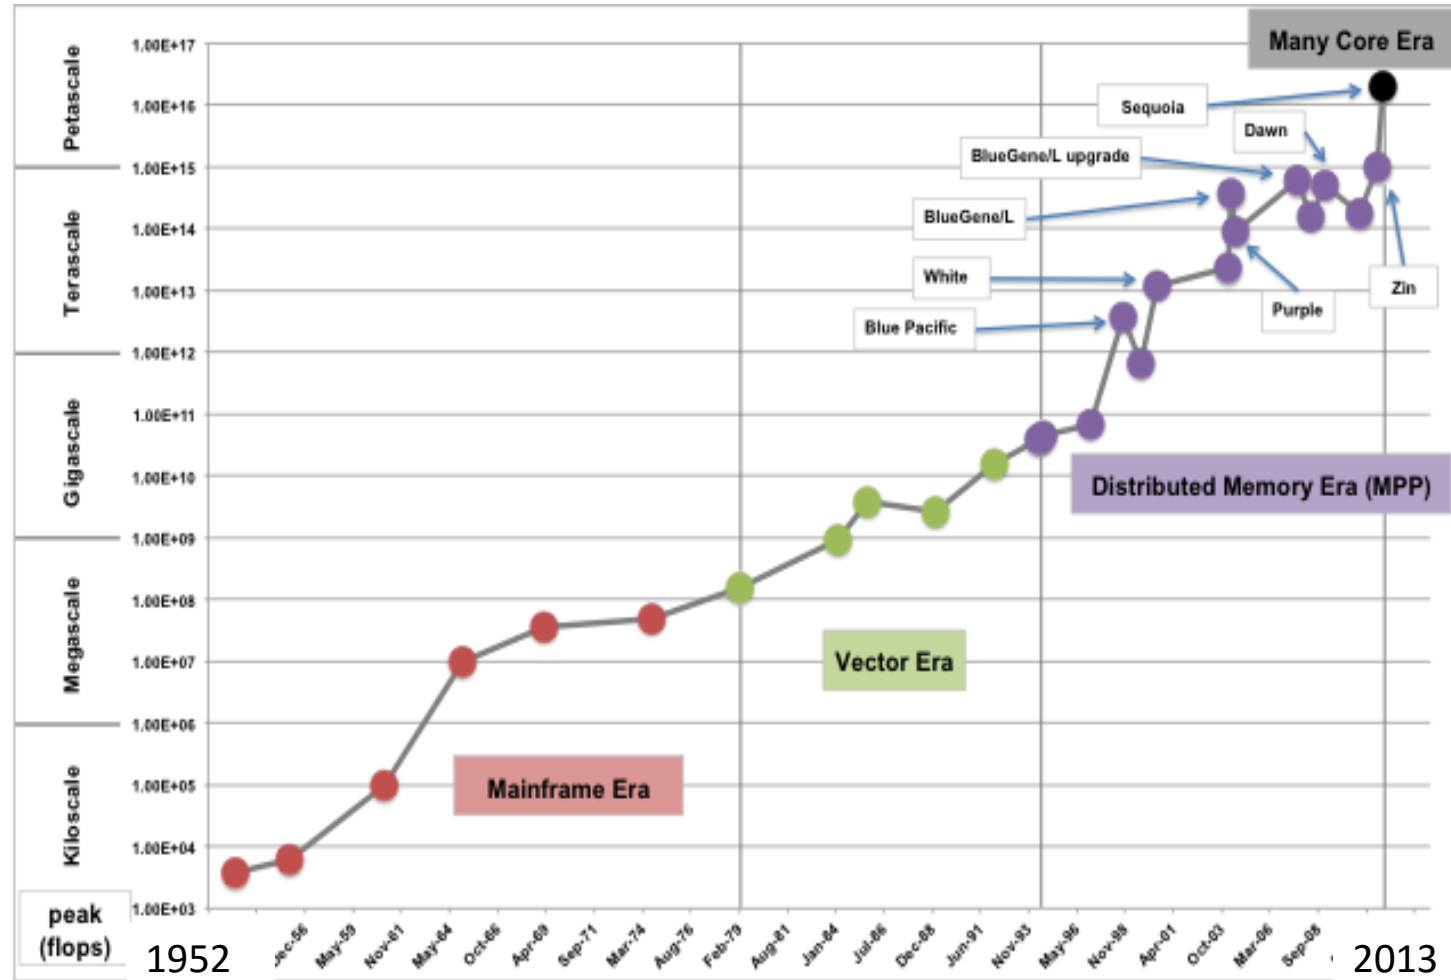
- *'Resiliency' against natural and man-made conditions*
- *Cross-sector dependencies*

Opportunity for broader societal impacts

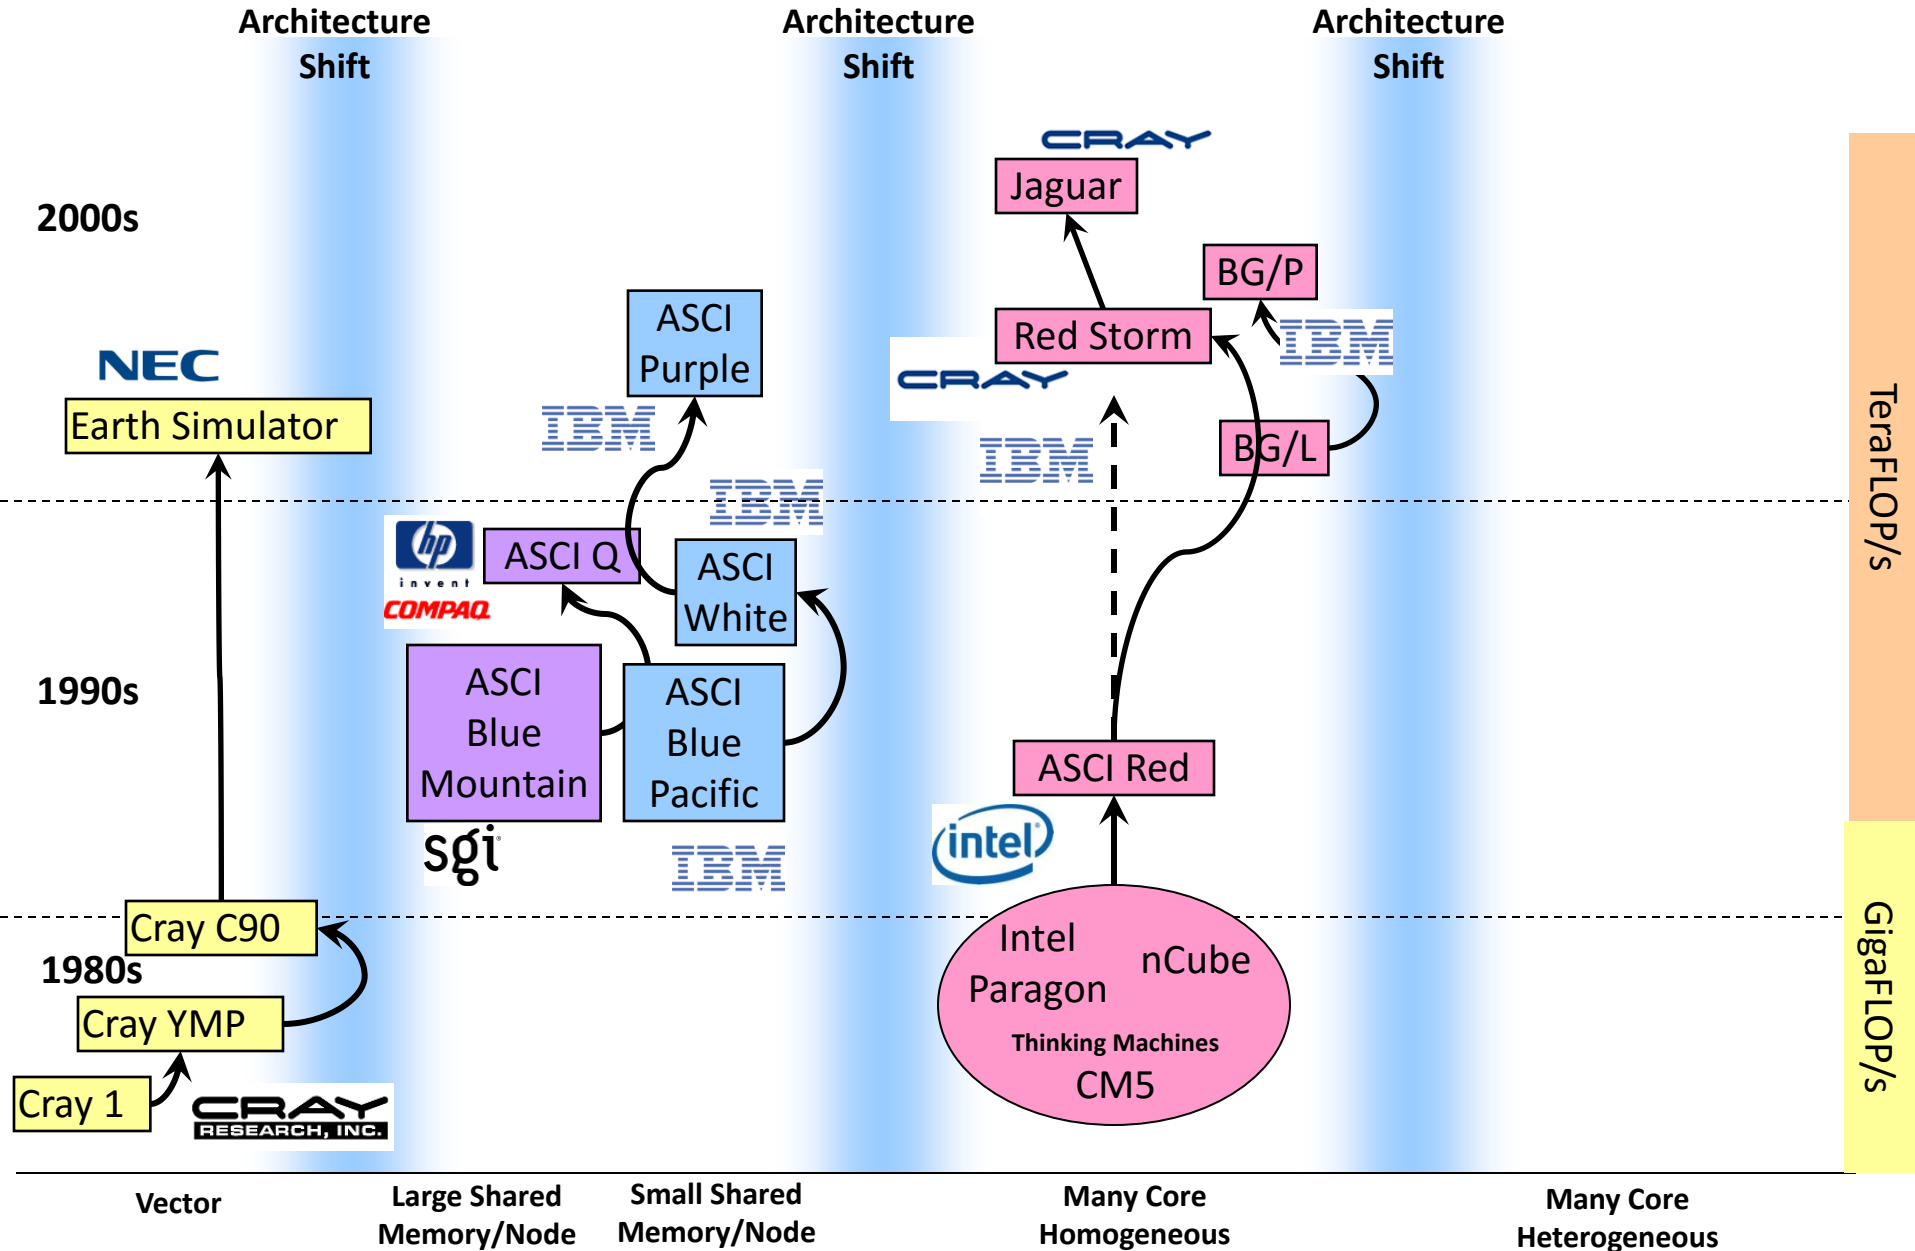
- *Extreme power requirements becomes a barrier to innovation.*

DOE has had a transformative and international leadership role in high-performance computing

Computer Performance



Family Tree of Recent Top Computing Architectures



Family Tree of Recent Top Computing Architectures

Architecture
Shift

Architecture
Shift

Architecture
Shift

2020s

2010s

2000s

100s PetaFLOP/s

10s PetaFLOP/s

Peta FLOP/s

~150 PF
IBM
IBM Power8
NVIDIA Volta

IBM

IBM

Summit

Sierra

27 PF
Cray XC30
Intel
Haswell/Knights
Landing
Intel Phi

CRAY

Trinity

CRAY

Titan

27 PF
Cray XK
AMD x86_64
NVIDIA K20x

20 PF
IBM Bluegene
PowerPC
No Coprocessor

IBM

Sequoia

IBM

IBM

Mira

10 PF
IBM Bluegene
PowerPC
No Coprocessor

Vulcan

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

IBM

DOE: Lead agency for NSCI

THE WHITE HOUSE

Office of the Press Secretary

For Immediate Release

July 29, 2015

EXECUTIVE ORDER

- - - - -

CREATING A NATIONAL STRATEGIC COMPUTING INITIATIVE

By the authority vested in me as President by the Constitution and the laws of the United States of America, and to maximize benefits of high-performance computing (HPC) research, development, and deployment, it is hereby ordered as follows:

Section 1. Policy. In order to maximize the benefits of HPC for economic competitiveness and scientific discovery, the United States Government must create a coordinated Federal strategy in HPC research, development, and deployment. Investment in HPC has contributed substantially to national economic prosperity and rapidly accelerated scientific discovery. Creating and deploying technology at the leading edge is vital to advancing my Administration's priorities and spurring innovation. Accordingly, this order establishes the National Strategic Computing Initiative. The Initiative shall be a whole-of-government, multi-agency strategy to coordinate and execute in collaboration with the private sector to maximize the benefits of HPC.

Over the past decade, HPC has been maintained and deployment of HPC performance on a par with the industry, and as coming decades will see increasing demand challenges and opportunities that require a cohesive Government and a private sectors.

It is the policy of the United States to enhance its scientific position in HPC through a coordinated Federal strategy.

(1) The United States shall support scientific research and development in HPC.

(2) The United States shall support collaborative government, industry, and academic research and development in HPC.

Sec. 2. Objectives. Executive departments, agencies, and offices (agencies) participating in the NSCI shall pursue five strategic objectives:

- (1) Accelerating delivery of a capable exascale computing system that integrates hardware and software capability to deliver approximately 100 times the performance of current 10 petaflop systems across a range of applications representing government needs.
- (2) Increasing coherence between the technology base used for modeling and simulation and that used for data analytic computing.
- (3) Establishing, over the next 15 years, a viable path forward for future HPC systems even after the limits of current semiconductor technology are reached (the "post-Moore's Law era").
- (4) Increasing the capacity and capability of an enduring national HPC ecosystem by employing a holistic approach that addresses relevant factors such as networking technology, workflow, downward scaling, foundational algorithms and software, accessibility, and workforce development.
- (5) Developing an enduring public-private collaboration to ensure that the benefits of the research and development advances are, to the greatest extent, shared between the United States Government and industrial and academic sectors.

Ecosystem:

1 W/ 1TF wearables

100TF wall-socket systems

1PF rack

...

The White House

Office of the Press Secretary

For Immediate Release

January 30, 2015

FACT SHEET: President Obama's Precision Medicine Initiative

Building on President Obama's announcement in his State of the Union Address, today the Administration is unveiling details about the Precision Medicine Initiative, a bold new research effort to revolutionize how we improve health and treat disease. Launched with a \$215 million investment in the President's 2016 Budget, the Precision Medicine Initiative will pioneer a new approach to medicine that promises to accelerate biomedical research, develop new tools, knowledge, and treatments that are best for which patients.

Most medical treatments have been developed as a result of this "one-size-fits-all" approach, which is successful for some patients but not for others. The emergence of precision medicine, which uses people's genes, environment, and lifestyle to develop clinicians tools to better understand and improve a patient's health, disease, and treatments will be most effective.

The White House

Office of the Press Secretary

For Immediate Release

January 28, 2016

Memorandum -- White House Cancer Moonshot Task Force

January 28, 2016

MEMORANDUM FOR THE HEADS OF EXECUTIVE DEPARTMENTS AND AGENCIES

SUBJECT: White House Cancer Moonshot Task Force

Cancer is a leading cause of death, and cancer incidence is expected to increase worldwide in the coming decades. But today, cancer research is on the cusp of major breakthroughs. It is of critical national importance that we accelerate progress towards prevention, treatment, and a cure -- to double the rate of progress in the fight against cancer -- and put ourselves on a path to achieve in just 5 years research and treatment gains that otherwise might take a decade or more. To that end, I hereby direct the following:

Section 1. White House Cancer Moonshot Task Force. There is established, within the Office of the Vice President, a White House Cancer Moonshot Task Force (Task Force), which will focus on making the most of Federal investments, targeted incentives, private sector efforts from industry and philanthropy, patient engagement initiatives, and other mechanisms to support cancer research and enable progress in treatment and care. The Vice President shall serve as Chair of the Task Force.

Section 2. Membership of the Task Force. In addition to the Vice President, the Task Force shall consist of the heads of the executive branch departments, agencies, and offices listed below:

- (i) the Department of Defense;
- (ii) the Department of Commerce;
- (iii) the Department of Health and Human Services;
- (iv) the Department of Energy;
- (v) the Department of Veterans Affairs;

DOE-NCI Opportunites

- DOE will build on its longstanding responsibilities in biomedical research - from radiation biology and associated cancer research questions relevant to nuclear security, to the genomics revolution essential in bioenergy research. The extraordinary tools in the DOE national laboratories that spawned and drove the genomics revolution, such as leading-edge computation, will now provide unique capabilities for addressing driving mission needs within the NCI.
- DOE can bring technology to the fight in a potentially disruptive way. In partnership with the NCI, DOE is poised to bring the next generation thinking to this class of challenge.
- These pilots are a natural driver for extreme computing. Cancer is a field replete with data, and in many cases progress is hindered by our natural abilities to learn everything we can from its full richness.
- DOE extreme computing is pushing the intersection of big data analytics, machine learning and modelling and simulation.

DOE, NCI and the Cancer Moonshot

- Machine learning builds on ideas in pattern recognition in images, numerical data, text, etc. It could be defined as statistical inference of patterns in data. In supervised learning, a computer is trained to look for specific signatures and works to identify those in data. But we are working to push towards ‘unsupervised learning’, where the system can develop its own sense of what data might contain. Medical image analysis is largely human limited. There is far more richness in the data than humans can apprehend or appreciate.
- Big data reflects the growing technological ability to capture, aggregate, and process an ever-greater volume, velocity, and variety of data. Big datasets are large, diverse, complex, longitudinal, and/or distributed datasets generated from increasingly ubiquitous and rich sensing, instrumentation, images and so on.
- Modeling and simulation is core to predicting complex behaviors of systems without actually testing them in real life – in many cases where important decisions need to be made and where testing is impossible (nuclear weapons, climate,...). This drives the DOE exascale computing effort. Predicting the health trajectories for patient outcomes of biological mechanisms present in cancer quickly becomes an exascale challenge.
- What matters is what you can do with these. Cancer challenges include analysis, capture, data curation, search, sharing, storage, transfer, visualization, querying and information privacy. Large data bases could benefit from novel ways to mine and exploit the information. The opportunity exists to leverage big data into data driven models, which complement the fundamental models that the experimental cancer community discovers.

DOE/NCI Pilots in Predictive Oncology

- The complexity of cancer is too much for the human to manage. It will take a bold step to merge machine learning, big data analytics and modeling & simulation together into a single framework to reveal cancer's weak-points. This is the core technology for the DOE/NCI pilots.
- The DOE/NCI **Joint Design of Advanced Computing Solutions for Cancer** joins DOE and cancer scientists together pushing both the frontiers of high-performance computing and predictive oncology.

Pilot 1 focuses on tailoring therapeutics and is a framework for predictive models for preclinical, individualized drug screening that will also drive the initial integration of machine learning functionality into the CORAL computer hardware.

Pilot 2 is aimed at developing a RAS(rat sarcoma)-complex interaction model through driving novel computational techniques in dynamic multi-scale simulations. It will push the development of machine learning and has the potential to drive understanding for prevention, detection and therapeutics.

Pilot 3 will inform both prevention and therapeutics and will create a modeling framework for predictive simulations of patient health trajectories which will drive the integration of big data analytics with data-driven modeling and simulation for the CORAL computer architectures.

CORAL as a transformative architecture



Exascale Computing Initiative

Timeline

- Trinity System FY16 40 pF
- CORAL Systems FY16-FY18 100-300pF
- APEX Systems FY19-FY20 < 1EF
- Exascale Systems FY22-FY23 1-10EF

DOE recognizes that leadership in precision medicine in general and cancer specifically will be advanced greatly by collaboration across disciplines to develop the next generation computational technologies that meld modeling and simulation, machine learning and big data analytics. This will be transformative and will also add to DOE's world leading computational toolbox for its energy, science and security missions.